TESTIMONY OF IRVING WLADAWSKY-BERGER
VICE PRESIDENT
TECHNOLOGY AND STRATEGY
IBM CORPORATION
BEFORE THE
U.S. HOUSE OF REPRESENTATIVES
COMMITTEE ON SCIENCE
MAY 13, 2004

Good morning, Mr. Chairman and members of the committee and thank you for inviting me to be with you today. My name is Irving Wladawsky-Berger. I am Vice President, Technology and Strategy at the IBM Corporation. I genuinely appreciate the opportunity to offer you our perspective on the questions before the committee.

Having been associated with high-performance computing for more than 30 years, I think it is important to share with you the fundamental shift we see happening in supercomputing and the role it will play in determining our nation's position in the global economy.

First, I'd like to thank Representative Biggert for her leadership on the important issue of high performance computing and express my appreciation to all of you for considering HR 4218 today. It is critical that our nation support the basic tenets of this bill to: 1) assure US researchers access to the most advanced high-performance computing systems available; 2) assure balanced progress on all aspects of high-performance computing; and 3) assure an adequate interagency planning process to maintain continued US leadership.

Second, I think a little historical perspective may be helpful.

There was a time when many in the US feared that we would lose leadership in this critical area to the Japanese IT industry. Instead, thanks to the combined efforts of industry, academia and government, the US established a strong leadership position in high-performance computing.

Why was this so important?

Because we needed supercomputing to address such grand challenges as:

- Enhancing military systems
- Building more energy-efficient cars and airplanes
- Designing better drugs
- Forecasting weather and predicting global climate change
- Improving environmental modeling, and
- Understanding the formation of galaxies, the nature of new materials, and the structure of biological molecules

Our leadership in high-performance computing technology allowed us to maintain our leadership internationally in these areas, and we did so with machines that are rudimentary by today's standards. Ten years ago, for example, the number one ranked machine on the world's Top 500 list of supercomputers performed 125 billion calculations per second. Today that computer would not even make the list.

I believe that supercomputing is even more important today than it was in the 1990s when we established our leadership. And if anything it is even more important now that we not only maintain but extend our leadership.

The same economic and social forces that are making PCs, the Internet, wireless and other technologies ubiquitous are transforming the high-performance computing segment.

Supercomputers have become so much less expensive and so much more powerful that they can now be applied in areas where they were never before affordable. In effect, the country's continuing commitment to this technology is making it possible to address *new* grand challenges. It is imperative that we do so.

- EPA, for example, will use a powerful new supercomputer to assess the risks to human health and the environment posed by exposure to chemical and air pollution and other agents.

- The State University of New York at Buffalo will use high-performance computers at our Deep Computing On Demand center to study human proteins and target drugs for cancer, Alzheimer's, AIDS, multiple sclerosis and other diseases.

Life sciences clearly represents an entirely new set of Grand Challenges for supercomputing with the potential to revolutionize health care in this country and the rest of the world. We cannot afford to ignore it.

We created our Blue Gene supercomputer initiative – ironically using the same chips found in game-players -- to tackle the Grand Challenge of protein folding. But there are other milestones we must reach -- including the simulation of drug interactions with human cells -- that are beyond today's systems. Today, we ultimately test new drugs on human beings. We know the cost and human suffering inherent in this process can be reduced dramatically over time with very sophisticated high-performance simulations leveraging many petaflops of computing power.

But supercomputing is no longer limited to the "classic" Grand Challenges.

At IBM we have described an emerging state of business called On Demand. This is fundamentally what happens when we become an information-based society with everything and everyone connected using open standards, and with computing power, storage and networking essentially unlimited.

Real-time applications and unprecedented amounts of data are creating an environment in which supercomputing is a requirement. Real-time transactions and data without real-time analysis and insight are no longer enough. We see this already in areas as diverse as fraud detection and customer relationship management. We believe supercomputing is rapidly becoming an essential part of the modern computing fabric.

Omnipresent communications keeps the world online and in touch 24-hours a day. Some experts believe that by 2006 the number of devices attached to the Internet – everything from PCs, smart phones and set-top boxes to RFID tags, home appliances and automobiles – will approach 10 billion; the number of users will approximate 1 billion; the number of online buyers a half-billion; and the total amount of commerce $5.5 trillion. Indeed, the price/performance ratio of microprocessors has made them so affordable that they can be integrated in huge numbers into everything from oil well drilling rigs and home appliances to vending machines and automobiles. Adidas is even putting them in running shoes.

Open standards are integrating all this technology and enabling it to amass and transmit information. The availability of information on such a scale and timeframe leads to decisions, decisions to actions, actions to change and change to the need for response. The pace of change will only accelerate and its magnitude will only increase with the constant proliferation and integration of technology.

Given that prospect, it is not surprising that in a recent IBM survey of 400 chief executive officers worldwide, the ability to respond to change was cited as a major need. Those CEOs were calling for the ability to take all that information created by customers and competitors and process it in real time. More and more, it is important to solve complex problems that are critical to competing in a global marketplace that demands the highest quality products offered at attractive prices with the best possible customer service.

Supercomputers are an excellent tool to collect and analyze data; simulate and model problems; and create real-time solutions. The power of supercomputers helps industry and the scientific community to innovate and create solutions faster and at less cost.

It is only with high-performance computing that we can hope to do the real-time information analysis that will enable us to respond faster and more effectively to the developing challenges and growing opportunities all our institutions face. Examples include gathering data to meet security challenges, developing everything from airplanes to health-related items, meeting customer needs, simulating drug reactions in the body, and tracking climate and weather to better understand the environmental challenges of the modern world.

Supercomputers can permit just about all of society's institutions – not just the research community -- to understand change better and to act with precision. But to make supercomputers more ubiquitous and increasingly helpful in a wide range of problems in business, health care, education, national security and every other aspect of society, those supercomputers must be affordable.

High-performance computing is crossing the boundary between the lab and the rest of society and is on the road to becoming a ubiquitous and conventional part of the IT infrastructure. As such, it should continue to be a driver of economic growth, a strategic tool for our scientific and business communities, and a strong pillar of our competitiveness in a changing, often turbulent, global marketplace.

The United States must ensure that it will have the high-performance computing assets needed in order to prosper in a constantly changing environment. Clearly, that requires aggressive research, performed at a level commensurate with the environment of change that we face, including the application of high-performance computing to produce real innovation.

We need to foster an environment of innovation much the way The High Performance Computing Act of 1991 and the Federal High Performance Computing and Communications (HPCC) program did when they gave scientists, engineers and industry leaders increased access to high-performance computers, thus building the user community and advancing science.

Innovation has always been the strong suit of the United States. Today, innovation remains the key to maintaining our ability to compete in a changing global economy where technology, science and education are becoming widespread among developing as well as developed nations. And the most advanced technologies – like supercomputing -- remain the key to innovation and competitive advantage.

Let me turn now to the specific questions posed by the Committee.

How does high-performance computing affect U.S. industrial competitiveness?

Supercomputing today is more important than ever, especially with the massive amounts of data we collect, analyze and use as well as the increasing complexity of our world. This is true given the competitive environment we live in, with constant growth in Asia and the European Union. This is equally true at the level of the individual firm, where customers have become far more demanding in terms of responsiveness, quality and price. US businesses recognize the value of high-end computing, and want the benefit of affordable access to these tools.

High-end computing has become the third node of science and engineering. By bridging theory and experimentation with computing and simulation, American industry is able to address some of the most complex, computationally intensive problems. Application areas extend from aircraft and automobile design to fusion reactor and accelerator design to materials science to petroleum exploration. High-end computing extends the amount of science and engineering that can be supported by available computational resources.

Supercomputing is the preferred tool of analysis for the sheer mass of available digital data created by advances in processing capability and inexpensive communications. New applications include the processing of streaming data, analysis of video and audio data, real-time security scanning and new areas such as information-based medicine.

Consider what two of our customers are already doing:

- Locus Pharmaceuticals is using supercomputing to develop novel small molecule therapeutics for viral diseases like AIDS.

- General Motors is installing the industry's fastest supercomputer based on our own POWER4 technology to promote greater global collaboration, improve validation testing and reduce product-development costs. They expect it to shorten some vehicles' time-to-market by as much as four years.

High-performance computing is making it possible to provide high-powered analytical capability to traditional commercial applications. In the past, such systems have been focused on the management of data and planning models. Today, we are adding real-time operational capabilities, permitting analysis of the data and response to changing external situations.

Delivering these capabilities sooner rather than later will be vital to US industry's ability to compete in an economic and regulatory environment that is changing and often uncertain. Fortunately, a new business model for delivering high-end computing to US industry is emerging, effectively widening the application base and reducing costs. Specifically, I am referring to the offering of supercomputing power to customers over the Internet, helping to free them from the fixed costs and management responsibility of owning a supercomputer. In this model, a business is able to avoid technological risk as well as the financial risk associated with supercomputer ownership. That is especially important for companies with short-term projects or those with variable needs for supercomputing power.

Let me reiterate that price/performance plays a major role in making supercomputing a prime tool for competitive advantage. In that regard, scalable systems based on common components make it possible to reach a large user base, help reduce the cost and risk of development, and support a wide range of applications. Cooperation among application and systems developers is key to achieving sustained performance improvements. This is true in both the business world and in the academic and scientific arenas. In universities, where individual investigators lead small research teams and are funded by research grants, a system's price is a major factor in determining which projects proceed and at what rate.

As you know, however, technology by itself is not enough. Our competitiveness will also depend on fostering a broad set of sophisticated skills to match the sophistication and capability of the technology. Our analysis indicates a growing need for many special skills like technical and scientific solutions architects, business transformation consultants, software engineers and application portfolio managers. Highly skilled personnel are critical to the success of the IT industry which in turn is necessary for the economy's competitiveness.

That is one reason IBM invests heavily in training and professional development. This year we will invest over $750 million to help our employees build skills, including more than $200 million for "hot" skills. $400 million (53 percent) will be spent in the US. This investment will ensure that our employees have the skills that customers need in today's highly competitive IT world.

Are current efforts on the part of the federal civilian science agencies in high-performance computing sufficient to assure US leadership in this area? What should agencies such as the National Science Foundation and the Department of Energy be doing that they are not already doing?

The current efforts of federal civilian agencies are a good start, but are not enough to meet present demands. This is why we support the bill under consideration and its objectives of: 1) assuring US researchers access to the most advanced high-performance computing systems available; 2) assuring balanced progress on all aspects of high-performance computing; and 3) assuring an adequate interagency planning process to maintain continued US leadership. I believe that these steps will help the US to advance high-performance computing and maintain our position of leadership.

That leadership is based on a many factors. They include: sustainability, meeting application needs, developing algorithms, enhancing skills and creating test beds and partnerships between government, industry and universities. By these measures, there is no question that the US continues to lead the world in high-performance computing.

However, to meet the challenges and complexity of the world today, supercomputing must both meet the "classic" Grand Challenges and become ubiquitous in the solution of a wide variety of problems. There must be a concerted effort to do the necessary research and to move even faster than before if we are to maintain our leadership. In the final analysis, it is the cumulative presence of a variety of leadership characteristics, including skills, technologies (both hardware and software), application development, training methodologies, research, development, engineering, and manufacturing capabilities that will advance high-performance computing. Agencies must focus on all of these components to ensure success.

My fundamental view is that the US should increase its application capability in a cost-effective manner. The roadmap developed to meet these needs must be based on commercially viable technologies that can be optimized for application-specific needs.

The government agencies must work with the research communities and the private sector to define supercomputing applications and technology solutions. The federal government should not attempt to dictate market trends and architectural paths for industry. Rather, the government as a partner with industry should specify its critical needs and work with industry to meet them. These partnerships are critical.

<u>Where are you targeting IBM's high-performance computing research efforts? Are there particular industrial sectors that will benefit in the near term from anticipated HPC developments?</u>

IBM’s research strategy revolves around solving complex scientific and business problems more quickly and at lower costs. We continue to aggressively evolve and improve our product line by developing advanced microprocessors which we then use to build scalable families of products. We are also conducting considerable research to overcome obstacles to high degrees of parallelism.

We are doing a number of things to advance our systems, such as:

- Studying cost effective, uniprocessor building blocks that take advantage of the ability to run multiple system activities at the same time (concurrency, i.e. interconnecting main memory, storage, various caches, and then processor execution units and algorithms and application software).

- Recognizing that sustained system performance is more than just hardware, but includes also application development performance and application execution performance.

- Bringing evolutionary technological improvements to current systems with functional integration at the chip package level to provide differentiation.

- Continuing to perform research into the most difficult problems in silicon semiconductor technology and performance.

- Exploring open standard software as a critical aspect of future research and performance.

Our strategy requires that we pursue application-driven design through partnerships with the national labs, universities and government agencies. We are working to satisfy a spectrum of customer performance and price needs, so naturally we maintain continued partnerships with the technical and scientific community. We are engaged in a number of studies to combine new processor architectures with innovative high-performance networks.

Our strategy is based on the following beliefs:

1. HPC systems and applications are crucial, since they will continue to drive advancement in the computer industry. It is not an issue of just technology and hardware. Advancement depends on servers, software, storage, communications and a business model for low-cost delivery of high-performance computing.

2. Petaflop performance will advance in response to the needs of the scientific community, and growing application complexity requires adaptable high-performance computing systems. It is critical to listen to users and then focus on and develop the applications that meet their needs.

3. Architecture should scale up and scale out. We have pioneered both these models. We are committed to sustainable models and long-term viability as well as to ensuring that our customers have the greatest performance for the least amount of money.

4. Simulation and modeling are key to solving 21st century problems.

5. Partnerships between government, universities and industry are critical.

Therefore, our research strategy involves working closely with the Federal government in general and not solely with the agencies within the jurisdiction of the Science Committee, like the Department of Energy's Office of Science and the National Science Foundation, but very actively with other agencies, such as the Department of Defense and the Department of Energy's Defense programs. In this regard, I believe that the National Institutes of Health should place greater focus on the power that supercomputing could provide for further advances in the life sciences.

We view each of our government collaborations as an opportunity to undertake Grand Challenge applications and address the most complex problems of our times. Our view is that we should leverage our systems expertise in these arenas. These partnerships are valuable to industry, universities and government and we all benefit in unique ways. For a company like IBM, for instance, these projects are relevant to our commercial business and we can leverage this opportunity for learning and importing these new ideas into our products.

Industrial sectors that will benefit include: the life sciences, aircraft and automotive manufacturers, pharmaceutical companies, petroleum companies, and consumer products businesses.

Conclusion

It is critical that high-performance computing in the United States advance to meet the challenges of our complex world. Meeting our applications needs, the needs of our scientists and our businesses, and the skill demands of the 21st century will help us to advance high-performance computing and keep the US at the keen edge of innovation.

HR 4218 will help us accomplish this goal. Its emphasis on a mix of leadership, partnerships, powerful and affordable systems, and a strong focus on basic research will keep the US competitive and help us maintain the innovative spirit that has made us global leaders in technology and the most prosperous society on earth.